

Streaming Graph Partitioning in the Planted Partition Model

Charalampos E. Tsourakakis

Brown University

`charalampos_tsourakakis@brown.edu`

August 21, 2014

Abstract

The sheer increase in the size of graph data has created a lot of interest into developing efficient distributed graph processing frameworks. Popular existing frameworks such as GRAPHLAB and PREGEL rely on balanced graph partitioning in order to minimize communication and achieve work balance.

In this work we contribute to the recent research line of streaming graph partitioning [Sta14, SK12, TGRV14] which computes an approximately balanced k -partitioning of the vertex set of a graph using a single pass over the graph stream using degree-based criteria. This graph partitioning framework is well tailored to processing large-scale and dynamic graphs. In this work we introduce the use of higher length walks for streaming graph partitioning and show that their use incurs a minor computational cost which can significantly improve the quality of the graph partition. We perform an average case analysis of our algorithm using the planted partition model [CK01, McS01]. We complement the recent results of Stanton [Sta14] by showing that our proposed method recovers the true partition with high probability even when the gap of the model tends to zero as the size of the graph grows. Furthermore, among the wide number of choices for the length of the walks we show that the proposed length is optimal. Finally, we conduct experiments which verify the value of the proposed method.

1 Introduction

The size of graph data that are required to be processed nowadays is massive. For instance, the Web graph amounts to at least one trillion of links [goo] and Facebook in 2012 reported more than 1 billion of users and 140 billion of friend connections. Furthermore, graphs of significantly greater size emerge by post-processing various other data such as image and text datasets. This sheer increase in the size of graphs has created a lot of interest in developing distributed graph processing systems [KTF09, LGK⁺10, MAB⁺10], in which the graph is distributed across multiple machines. A key problem towards enabling efficient graph computations in such systems is the **NP**-hard problem of *balanced graph partitioning*. High-quality partitions ensure low volume of communication and work balance.

Recently, Stanton and Kliot [SK12] introduced a streaming graph partitioning model. This line of research despite being recent and lacking theoretical understanding has already attracted a lot of interest. Several existing systems have incorporated this model such as PowerGraph [GLG⁺12]. The framework of FENNEL has been adapted by PowerLyra [CSC⁺13], which has been included into the most recent GRAPHLAB version [bic, LGK⁺10] yielding significant speedups for various iterative computations. Stanton performed an average case analysis of two streaming algorithms, and explained their efficiency despite the pessimistic worst case analysis [Sta14]. Despite the fact that existing established heuristics such as METIS typically outperform streaming algorithms, the latter are well tailored to today's needs for processing dynamic graphs and big graph data which do not fit in the main memory. They are computationally cheap and provide partitions of comparable quality. For instance, FENNEL on the Twitter network with more than 1.4 billion edges performed comparably well with METIS for a wide variety of settings, requiring 40 minutes of running time, whereas METIS $8\frac{1}{2}$ hours.

So far, the work on streaming graph partitioning is based on computing the degrees of incoming vertices towards each of the k available machines. Equivalently, this can be seen as performing a one step random walk from the incoming vertex. A natural idea which is used extensively in the literature of graph partitioning [LS90, ST08, ZW04] is the use of higher length walks. In this work we introduce this idea in the setting of streaming graph partitioning. At the same time we maintain the time efficiency of streaming graph partitioning algorithms which make them attractive to various graph processing systems [GLG⁺12, KTF09, LGK⁺10, MAB⁺10].

Summary of our contributions. Our contributions can be summarized as follows:

- Our proposed algorithm introduces the idea of using higher length walks for streaming graph partitioning. It incurs a negligible computational cost and significantly improves the quality of the partition. We perform an average case analysis on the planted partition model, c.f. Section 3.2 for the description of the model. We complement the recent results of [Sta14] which require that the gap $p - q$ of the planted partition model is constant in order to recover the partition *whp*¹, by allowing gaps $p - q$ which asymptotically tend to 0 as n grows.
- Among the wide number of choices for the length of t -walks where $t \geq 2$, we show that walks of length 2 are optimal as they allow the smallest possible gap for which we can guarantee recovery *whp*, c.f. Section 3.5.
- We evaluate our method on the planted partition model and we provide a machine learning application which illustrates the benefits of our method.

Structure of the paper. The remainder of the paper is organized as follows. Section 2 discusses briefly related work. Section 3 presents the proposed algorithm and its analysis. Section 4 presents the experimental evaluation. Section 5 concludes.

2 Related Work

Section 2.1 briefly reviews existing work. Section 2.2 states the probabilistic tools we use in Section 3. A complete account can be found in [DP09, MU05].

2.1 Existing Work

Balanced graph partitioning. The *balanced graph partitioning* problem is a classic NP-hard problem of fundamental importance to parallel and distributed computing. The input to this problem is an undirected graph $G(V, E)$ and an integer $k \in \mathbb{Z}^+$, the output is a partition of the vertex set in k balanced sets such that the number of edges across the clusters is minimized. We refer to the k sets as *clusters* or *machines* interchangeably. Formally, the balance constraint is defined by the imbalance parameter ν . Specifically, the

¹An event \mathcal{E}_n occurs *with high probability*, or *whp* for brevity, if $\lim_{n \rightarrow \infty} \Pr[\mathcal{E}_n] = 1$.

(k, ν) -balanced graph partitioning asks to divide the vertices of a graph in k clusters each of size at most $\nu \frac{n}{k}$, where n is the number of vertices in G . The case $k = 2, \nu = 1$ is equivalent to the **NP**-hard minimum bisection problem. Several approximation algorithms, e.g., [FK02], and heuristics, e.g., [KL70], exist for this problem. When $\nu = 1 + \epsilon$ for any desired but fixed $\epsilon > 0$ there exists a $O(\epsilon^{-2} \log^{1.5} n)$ approximation algorithm [KR04]. When $\nu = 2$ there exists an $O(\sqrt{\log k \log n})$ approximation algorithm based on semidefinite programming (SDP) [KNS09]. Due to the practical importance of k -partitioning there exist several heuristics, among which METIS [SKK00] stands out for its good performance. Survey [BMS⁺] summarizes many popular existing methods for the balanced graph partitioning problem.

Streaming balanced graph partitioning. Despite the large amount of work on the balanced graph partitioning problem, neither state-of-the-art approximation algorithms nor heuristics such as METIS are well tailored to the computational restrictions that the size of today’s graphs impose.

Stanton and Kliot introduced the streaming balanced graph partitioning problem, where the graph arrives in the stream and decisions about the partition need to be taken with on the fly quickly [SK12]. Specifically, the vertices of the graph arrive in a stream with the set of edges incident to them. When a vertex arrives, a partitioner decides where to place the vertex “on the fly”, using limited computational resources (time and space). A vertex is never relocated after it becomes assigned to one of the k machines. In [TGRV14] well-performing decision strategies for the partitioner were introduced inspired by a generalization of optimal quasi-cliques [TBG⁺13] to k -partitioning. Both [SK12] and [TGRV14] can be adapted to edge streams. Nishimura and Ugander [NU13] consider a realistic variation of the original problem where the constraint of a single pass over the graph stream is relaxed to allowing multiple passes. Stanton showed that streaming graph partitioning algorithms with a single pass under an adversarial stream order cannot approximate the optimal cut size within $o(n)$. The same claim holds also for random stream orders [Sta14]. Finally, Stanton [Sta14] analyzes two variants of well performing algorithms from [SK12] on random graphs. Specifically, Stanton proves that if the graph G is sampled according to the planted partition model, then the two algorithms despite their similarity can perform differently and that one of the two can recover the true partition *whp*, assuming that inter-, intra-cluster edge probabilities are constant, and their gap is a large constant.

Planted partition model. Jerrum and Sorkin [Jer98] studied the planted bisection model, a random undirected graph with an even number of vertices. According to this model, each half of the vertices is assigned to one of two clusters. Then, the probability of an edge (i, j) is p if i, j have the same color, otherwise $q < p$. We will refer to p, q as the intra- and inter-cluster probabilities. Their difference $p - q$ will be referred to as the *gap* of the model. Condon and Karp [CK01] studied the generalization of the planted bisection problem where instead of having only two clusters, there exist k clusters of equal size. The probability of an edge is the same as in the planted bisection problem. They show that the hidden partition can be recovered *whp* if the gap satisfies $p - q \geq n^{-1/2+\epsilon}$. McSherry [McS01] presents a spectral algorithm that recovers the hidden partition in a random graph *whp* if $p - q = \Omega(n^{-1/2+\epsilon})$. Recently, Van Vu gave an alternative algorithm [Vu14] to obtain McSherry's result. Zhou and Woodruff [ZW04] showed that if $p = \Theta(1), q = \Theta(1), p - q = \Omega(n^{-1/4})$, then a simple algorithm based on squaring the adjacency matrix of the graph recovers the hidden partition *whp*.

Random walks. The idea of using walks of length greater than one [ZW04] is common in the general setting of graph partitioning. Lovász and Simonovits [LS90] show that random walks of length $O(\frac{1}{\phi})$ can be used to compute a cut with sparsity at most $\tilde{O}(\sqrt{\phi})$ if the sparsest cut has conductance ϕ . Later, Spielman and Teng [ST08] provided a local graph partitioning algorithm which implements efficiently the Lovász-Simonovits idea.

2.2 Theoretical Preliminaries

A useful lemma that we use extensively is Boole's inequality, also known as the union bound.

Lemma 1 (Union bound) *Let A_1, \dots, A_n be events in a probability space Ω , then*

$$\Pr[\cup_{i=1}^n A_i] \leq \sum_{i=1}^n \Pr[A_i].$$

The following theorem is due to Markov and its use is known as the first moment method.

Theorem 1 (First Moment Method) *Let X be a non-negative random variable with finite expected value $\mathbb{E}[X]$. Then, for any real number $t > 0$,*

$$\Pr[X \geq t] \leq \frac{\mathbb{E}[X]}{t}.$$

The following theorem is due to Chebyshev and its use is known as the second moment method.

Theorem 2 (Second Moment Method) *Let X be a random variable with finite expected value $\mathbb{E}[X]$ and finite non-zero variance $\text{Var}[X]$. Then, for any real number $t > 0$,*

$$\Pr[|X - \mathbb{E}[X]| \geq t] \leq \frac{\text{Var}[X]}{t^2}.$$

Finally, we use the following Chernoff bounds for independent and negatively correlated random variables.

Theorem 3 (Multiplicative Chernoff Bound) *Let $X = \sum_{i=1}^n X_i$ where X_1, \dots, X_n are independent random variables taking values in $[0, 1]$. Also, let $\delta \in [0, 1]$. Then,*

$$\Pr(X \leq (1 - \delta)E(X)) \leq e^{-\delta^2 E(X)/2}.$$

Definition 1 (Negatively correlated random variables) *Let X_1, \dots, X_n be random binary variables. We say that they are negatively correlated if and only if for all sets $I \subseteq [n]$ the following inequalities are true:*

$$\Pr(\forall i \in I : X_i = 0) \leq \prod_{i \in I} \Pr(X_i = 0)$$

and

$$\Pr(\forall i \in I : X_i = 1) \leq \prod_{i \in I} \Pr(X_i = 1).$$

Theorem 4 *Let $X = \sum_{i=1}^n X_i$ where X_1, \dots, X_n are negatively correlated binary random variables. Also, let $\delta \in [0, 1]$. Then,*

$$\Pr(X \leq (1 - \delta)E(X)) \leq e^{-\delta^2 E(X)/2}.$$

3 Proposed Algorithm

Section 3.1 introduces our notation and Section 3.2 presents in detail the random graph model we analyze. Section 3.3 provides two useful lemmas used in Sections 3.4 and 3.5. Section 3.4 shows an efficient way to recover the true partition of the planted partition model *whp* using walks of length $t = 2$ even when $p - q = o(1)$, i.e., the gap is asymptotically equal to 0. Section 3.5 shows that when the length of the walk t is set to 2, then we obtain the smallest possible *gap* $p - q$ for which we can guarantee recovery of the partition *whp*. We do not try to optimize constants in our proofs, since we are interested in asymptotics. Our proofs use the elementary inequalities $1 - p \leq e^{-p}$, $\binom{n}{k} \leq \left(\frac{en}{k}\right)^k$ and the probabilistic tools presented in Section 2.2.

3.1 Notation

Let $G(V, E)$ be a simple undirected graph, where $|V| = n, |E| = m$. We define a *vertex partition* $\mathcal{P} = (S_1, \dots, S_k)$ as a family of pairwise disjoint vertex sets whose union is V . Each set S_i is assigned to one of k machines, $i = 1, \dots, k$. We refer to each S_i as a *cluster* or *machine*. Throughout this work, we assume $k = \Theta(1)$. Let $\partial e(\mathcal{P})$ be the set of edges that cross partition boundaries, i.e., $\partial e(\mathcal{P}) = \cup_{i=1}^k e(S_i, V \setminus S_i)$. The fraction of edges cut λ is defined as $\lambda = \frac{|\partial e(\mathcal{P})|}{m}$. The imbalance factor or normalized maximum load ρ is defined as $\rho = \max_{1 \leq i \leq k} \frac{|S_i|}{\frac{n}{k}}$. Notice that ρ satisfies the double inequality $1 \leq \rho \leq k$. When $\rho = 1$ we have a perfectly balanced partition. At the other extreme, when $\rho = k$ all the vertices are placed in one cluster, leaving $k - 1$ clusters empty. In practice, there is a constraint $\rho \leq \nu$ where ν is a value imposed by application-specific restrictions. Typically, ν is close to 1. We omit floor and ceiling notation for simplicity, our results remain valid.

3.2 Planted Partition Model

The model $G(n, \Psi, P)$ is a generalization of the classic Erdős-Rényi graph [Bol01]. The first parameter n is the number of vertices. The second parameter of the model is the function $\Psi : [n] \rightarrow [k]$ which maps each vertex to one of k clusters. We refer to $C_i = \Psi^{-1}(i)$ as the i -th cluster, $i = 1, \dots, k$. The third parameter P is a $k \times k$ matrix such that $0 \leq P_{ij} \leq 1$ for all $i, j = 1, \dots, k$ which specifies the probability distribution over the edge

set. Specifically, a graph $G \sim G(n, \Psi, P)$ is generated by adding an edge with probability $P(\Psi(u), \Psi(v))$ between each pair of vertices $\{u, v\}$. Notice, that when all the entries of P are equal to p , then $G(n, \Psi, P)$ is equivalent to the $G(n, p)$ model. In this work we are interested into graphs that exhibit clustering. The planted partition model we analyze is the same as in [CK01, McS01, Sta14, ZW04]: $P_{ij} = p\mathbf{1}(i = j) + q\mathbf{1}(i \neq j)$ for all $i, j = 1, \dots, k$. We assume that $p > q = \Theta(1)$. For simplicity, we refer to this version of $G(n, \Psi, P)$ as $G(n, k, p, q)$.

3.3 Results and useful lemmas

We prove two simple lemmas that we use in the analysis of our algorithms. We refer to the vertex that arrives exactly after $i-1$ vertices as the i -th vertex. The first lemma states that after a super-logarithmic number of vertices, we have seen a fraction of $\frac{1}{2k}$ vertices from each cluster of the planted partition *whp*.

Lemma 2 *For all $i \geq D \log^{1+\delta}(n)$ where $D, \delta > 0$ are any positive constants, there exist at least $\frac{i}{2k}$ vertices that have already arrived from each cluster j whp, $j = 1, \dots, k$.*

Proof 1 *Fix any index $i \geq D \log^{1+\delta}(n)$. Define for each $j \in [k]$. the bad event \mathcal{A}_j that cluster j has less than $\frac{i}{2k}$ vertices. after the arrival of i vertices. Given our assumption on the random order of the stream, the first i vertices form a random i -subset of $[n]$. Let Y_j be the number of vertices from cluster j among the first i vertices. We observe that the distribution of Y_j is the hypergeometric distribution $H(n, n/k, i)$ ². Therefore, we obtain the following exact expression for the probability of the bad event \mathcal{A}_j :*

$$\Pr[\mathcal{A}_j] = \sum_{r=0}^{i/2k} \Pr[Y_j = r] = \sum_{r=0}^{i/2k} \frac{\binom{n/k}{r} \binom{n(k-1)/k}{i-r}}{\binom{n}{i}}.$$

Even if an asymptotic analysis using Stirling's formula is possible, a less tedious approach is possible. We express Y_j as the sum of i indicator variables $Y_{j,1}, \dots, Y_{j,i}$ where $Y_{j,l} = 1$ if and only if the l -th vertex v has $\Psi(v) = j$,

²Recall, the hypergeometric distribution $H(N, R, s)$ is a discrete probability distribution that describes the probability of choosing r red balls in s draws of balls without replacement from a finite population of size N containing exactly R red balls.

for all $l = 1, \dots, i$. Clearly, $\mathbb{E}[Y_j] = \frac{i}{k}$. Notice that even if the indicator random variables are not independent, they are negatively correlated. Therefore, Theorem 4 applies, obtaining

$$\Pr \left[Y_j \leq \left(1 - \frac{1}{2}\right) \frac{i}{k} \right] \leq e^{-\frac{i}{8k}} \leq e^{-\frac{D \log^{1+\delta}(n)}{8k}} \ll o(n^{-1}).$$

The proof is completed by taking a union bound over k machines and $(n - D \log^{1+\delta}(n))$ vertices. Specifically, let \mathcal{E} be the event that there exists an index $i \geq D \log^{1+\delta}(n)$ such that $\cup_{j=1}^k \mathcal{A}_j$ is true. Then,

$$\Pr[\mathcal{E}] \leq (n - D \log^{1+\delta}(n)) k o(n^{-1}) = o(1).$$

The second lemma states that as soon as $\log^{6+\delta}(n)$ vertices have arrived in the stream for any $\delta > 0$, then the remaining vertices have at least $\log^6(n)$ neighbors in the set of already arrived vertices.

Lemma 3 *Let $\delta > 0$ be any positive constant. After $\log^{6+\delta}(n)$ vertices have arrived, all remaining vertices in the stream have $\log^6(n)$ neighbors which reside in the k machines.*

Proof 2 *Let $i = g(n) \geq \log^{6+\delta}(n)$, be the number of vertices that have arrived in the incidence stream. Let \mathcal{E} be the event that there exists a vertex v that has arrived after the first i vertices with less than $\log^6(n)$ neighbors in the set of arrived vertices. We obtain that $\Pr[\mathcal{E}] = o(1)$ using the following upper bound.*

$$\begin{aligned} \Pr[\mathcal{E}] &\leq n \sum_{j=0}^{\log^6(n)} \binom{g(n)}{j} \left(\frac{p}{1-q}\right)^j (1-q)^{g(n)} \leq n \sum_{j=0}^{\log^6(n)} \left(\frac{g(n)e}{j}\right)^j \left(\frac{p}{1-q}\right)^j (1-q)^{g(n)} \\ &\leq Cn \log^6(n) \left(\frac{pg(n)e}{j(1-q)}\right)^{\log^6(n)} e^{-qg(n)} = o(1). \end{aligned}$$

3.4 Path-2 classification

The main theoretical result of this Section is that we can use paths of length 2 to recover the partition Ψ whp even when $p - q = o(1)$. Our algorithm avoids making final decisions for the first B vertices until the end of the process and uses the fact that vertices from the same cluster have more

common neighbors compared to a pair of vertices from different clusters. Our algorithm is shown below.

1. Place the first $B = \log^{6+\delta}(n)$ vertices in any of the k machines, marked as non-classified. Here, $\delta > 0$ is any positive constant.
2. Let S be a random sample of size $3k \log n$ vertices from the set of B non-classified vertices.
3. Let R be a random sample of $\log^6 n$ vertices from the set of B non-classified vertices.
4. For the j -th vertex, $B + 1 \leq j \leq n$, do the following:
 - For each $x \in S$ compute the number of common neighbors of j, x in R .
 - Let $M = (p^2 + (k - 1)q^2) \log^6 n$. Assign j to the same cluster with a vertex $x^* \in S$ which has at least $M - M^{2/3}$ common neighbors with j . Ties are always assigned to the vertex with the smallest id. Remove non-classified tag from x^* .
5. Perform the same procedure for the remaining, if any, non-classified vertices.

We prove the correctness of the algorithm. The next lemma states that when we obtain the random sample S , there always exists at least one vertex from each cluster of the partition. This is critical since our algorithm assigns each incoming vertex v to a representative vertex from cluster $\Psi(v)$. Among the various possible choices for a representative of cluster c , we choose the vertex u with the minimum vertex id, namely $u = \arg \min\{u \in S : \Psi(u) = c\}$.

Lemma 4 *Let S be a random sample of size $3k \log n$ vertices from a population of $j \geq \log^{6+\delta}(n)$ vertices. Then, with high probability there exists at least one representative vertex from each cluster of the planted partition in S*

Proof 3 *First, notice that by Lemma 2 there exist at least $j/2k \geq \frac{\log^{6+\delta}(n)}{2k}$ vertices from each cluster $i = 1, \dots, k$. Let \mathcal{E}_i be the event of failing to sample at least one vertex from cluster i of the partition, $i = 1, \dots, k$. We can upper bound the probability of the union $\cup_{i=1}^k \mathcal{E}_i$ of these bad events as follows:*

$$\begin{aligned}\Pr\left[\bigcup_{i=1}^k \mathcal{E}_i\right] &\leq k\left(1 - \frac{j}{2k}\right)^{3k \log n} \\ &\leq ke^{-3k \log n / 2k} = o(n^{-1}).\end{aligned}$$

The next theorem is our main theoretical result. It states that the algorithm recovers the true partition Ψ *whp*.

Lemma 5 *If $p = \Theta(1), q = \Theta(1), p - q = \omega\left(\frac{1}{\log(n)}\right)$, then all vertices are classified correctly whp. The algorithm runs in sublinear time.*

Proof 4 Correctness: *Let j be the index of the incoming vertex, $x \in S$. We condition on the event \mathcal{E} that there exists at least one vertex from each cluster of the planted partition in the sample S . Define Y_x as the number of triples (j, u, x) where $u \in R$ and $x \in S$ such that $\Psi(j) = \Psi(x)$. Similarly, let Z_x be the number of triples (j, u, x) , where $u \in R$ and $x \in S$ such that $\Psi(j) \neq \Psi(x)$. The expected values of Y_x, Z_x are respectively*

$$\mathbb{E}[Y_x] = (p^2 + (k-1)q^2) \log^6 n,$$

and

$$\mathbb{E}[Z_x] = (2pq + (k-2)q^2) \log^6 n.$$

A direct application of the multiplicative Chernoff bound, see Theorem 3, yields

$$\Pr\left[Y_x \leq \mathbb{E}[Y_x] - \mathbb{E}[Y_x]^{2/3}\right] \leq e^{-O(\mathbb{E}[Y_x]^{1/3})} = n^{-O(\log n)}$$

and

$$\Pr\left[Z_x \geq \mathbb{E}[Z_x] + \mathbb{E}[Z_x]^{2/3}\right] \leq e^{-O(\mathbb{E}[Z_x]^{1/3})} = n^{-O(\log n)}$$

Furthermore, due to our assumption on $p, q = \Theta(1)$ and the gap $p - q \gg \frac{1}{\log n}$ we obtain

$$(p - q)^2 \log^6 n \gg \left((p^2 + (k - 1)q^2)^{2/3} + (2pq + (k - 2)q^2)^{2/3} \right) \log^4 n \Leftrightarrow \\ \mathbb{E}[Y_x] - \mathbb{E}[Z_x] \gg \mathbb{E}[Y_x]^{2/3} + \mathbb{E}[Z_x]^{2/3} \Leftrightarrow \mathbb{E}[Z_x] + \mathbb{E}[Z_x]^{2/3} \ll \mathbb{E}[Y_x] - \mathbb{E}[Y_x]^{2/3}.$$

The above inequalities suggest that the number of 2-paths between j and any vertex $x \in S$ such that $\Psi(x) = \Psi(j)$ is significantly larger compared to the respective count between j, x' such that $\Psi(x') \neq \Psi(j)$ whp. Let \mathcal{B}_j be the event that j is misclassified. Combining the above inequalities with Lemma 4 results in

$$\Pr[\mathcal{B}_j] \leq \Pr[\bar{\mathcal{E}}] + \Pr[\mathcal{B}_j | \mathcal{E}] \Pr[\mathcal{E}] = o(n^{-1}).$$

By a union bound over $O(n)$ vertices, the proof is complete.

Running time: The algorithm can be implemented in $O(n \text{polylog}(n))$ time in expectation. Sampling $O(\text{polylog}(n))$ samples can be implemented in expected $O(\text{polylog}(n))$ time, c.f. [Knu07]. Also, checking whether a neighbor of an incoming vertex resides in a given machine can be done in $O(1)$ time by using hash tables to store the information within each machine. Notice that the number of edges in G is $O(n^2)$ and therefore the proposed algorithm is a sublinear time algorithm.

3.5 Path- t classification

We conclude this section by discussing the effect of the length of the walk $t \geq 2$. Intuitively, t should not be too large, otherwise the random walk will mix. We argue, that among all constant lengths $t \geq 2$, the choice $t = 2$ allows the smallest possible gap for which we can find the true partition Ψ . It is worth outlining that $t = 2$ is also in favor of the graph partitioning efficiency as well, since the smaller t is, the less operations are required. Our results extend the results of Zhou and Woodruff [ZW04] to the streaming setting. Our main theoretical result is the following theorem.

Theorem 5 *Let $t \geq 2, t = \Theta(1)$ be the length of a walk. If $p, q = \Theta(1)$ such that $p(1 - p), q(1 - q) = \Theta(1)$, then $t = 2$ results in the largest possible gap $p - q$ for which we can decide whether $\Psi(u) = \Psi(v)$ or not whp, where $u \neq v \in [n]$.*

To prove Theorem 5, we compute first the expected number of walks of length t between any two vertices in $G(n, k, p, q)$. Because of the special structure of the graph, we are able to derive an exact formula³.

Lemma 6 *Let $G \sim G(nk, k, p, q)$ and $p_t = A_{uv}^t, q_t = A_{uw}^t$ be the two types of entries that appear in A^t depending on whether $\Psi(u) = \Psi(v)$ and $\Psi(u) \neq \Psi(w)$ respectively. Then*

$$p_t = (k-1) \frac{n^{t-1}(p-q)^t}{k} + \frac{n^{t-1}(p+(k-1)q)^t}{k},$$

and

$$q_t = -\frac{n^{t-1}(p-q)^t}{k} + \frac{n^{t-1}(p+(k-1)q)^t}{k}.$$

Proof 5 *Let A be the (p, q) -adjacency matrix defined as $A_{uv} = p$ if $\Psi(u) = \Psi(v)$ and $A_{uv} = q$ if $\Psi(u) \neq \Psi(w)$ for each $u \neq v \in [n]$. It is easy to check that for any $t \geq 1$ the block structure of the planted partition is preserved. This implies that for any t , matrix A^t has the same block structure as A and therefore there are two types of entries in each row. Let p_t, q_t be these two types of entries in A^t . For $t = 1$, let $p_1 = p$ and $q_1 = q$. Then, by considering the multiplication of the u -th line of A^t with the v -th column of A we obtain*

$$p_{t+1} = np_t + (k-1)nqq_t,$$

and similarly by considering the multiplication of the u -th line of A^t with the w -th column of A we obtain

$$q_{t+1} = qnp_t + pnq_t + \frac{k-2}{k}nqq_t.$$

We can write the recurrence in a matrix form.

$$\begin{bmatrix} p_{t+1} \\ q_{t+1} \end{bmatrix} = M \times \begin{bmatrix} p_t \\ q_t \end{bmatrix}$$

where $M = \begin{pmatrix} pn & (k-1)nq \\ qn & pn + \frac{k-2}{k}nq \end{pmatrix}$. By looking the eigendecomposition of M , despite the fact that it is not symmetric, we can diagonalize it as

³Despite the large amount of work on the planted partition, we were not able to find a closed formula but only bounds in the existing literature.

$$M = USU^{-1},$$

where

$$S = \begin{pmatrix} n(p-q) & 0 \\ 0 & n(p+(k-1)q) \end{pmatrix}$$

and

$$U = \begin{pmatrix} -(k-1) & 1 \\ 1 & 1 \end{pmatrix}.$$

Given the fact

$$M^k = U \begin{pmatrix} (n(p-q))^k & 0 \\ 0 & (n(p+(k-1)q))^k \end{pmatrix} U^{-1}$$

and simple algebraic manipulations (omitted) we obtain that

$$p_t = (k-1) \frac{n^{t-1}(p-q)^t}{k} + \frac{n^{t-1}(p+(k-1)q)^t}{k},$$

and

$$q_t = -\frac{n^{t-1}(p-q)^t}{k} + \frac{n^{t-1}(p+(k-1)q)^t}{k}.$$

Now, we are able to prove Theorem 5.

Proof 6 (Theorem 5) Let $p_t = A_{uv}^t, q_t = A_{uw}^t$ where $u, v, w \in V(G)$ such that $\Psi(u) = \Psi(v) \neq \Psi(w)$ and A is defined as in Lemma 6. Also, define \bar{A} to be the result of the randomized rounding of A . By Lemma 6 we obtain

$$p_t - q_t = \frac{n^{t-1}(p-q)^t}{k^{t-1}}.$$

Notice that we substituted n by n/k as we G has n vertices, with exactly n/k vertices per cluster. Now, suppose $|A_{uv}^t - \bar{A}_{uv}^t| \leq \gamma$, where $\gamma > 0$ is large enough such that the inequality holds whp and will be decided in the following. Then, if $\Psi(j_1) = \Psi(j_2) = \Psi(u)$ we obtain the following upper bound

$$|\bar{A}_{uj_1}^t - \bar{A}_{uj_2}^t| \leq |A_{uj_1}^t - \bar{A}_{uj_1}^t| + |A_{uj_2}^t - \bar{A}_{uj_2}^t| = 2\gamma.$$

On the other hand if $\Psi(u) = \Psi(j_1) \neq \Psi(j_2)$, given that $|x| = |x + y - y| \leq |y| + |x - y| \rightarrow |x - y| \geq |x| - |y|$, we obtain the following lower bound

$$|\bar{A}_{uj_1}^t - \bar{A}_{uj_2}^t| \geq |A_{uj_1}^t - A_{uj_2}^t| - 2\gamma = \frac{n^{t-1}(p-q)^t}{k^t} - 2\gamma.$$

Therefore, if $\frac{n^{t-1}(p-q)^t}{k^{t-1}} - 2\gamma > 2\gamma$, then there exists a signal that allows us to classify the vertex correctly. In order to find γ we need to upper-bound the expectation of the non-negative random variable $Z = (\bar{A}_{uv}^t - A_{uv}^t)^2$.

Zhou and Woodruff prove $\mathbb{E}[Z] = \Theta(n^{2t-3})$, c.f. Lemma 5, [ZW04]. The proof of this claim is based on algebraic manipulation. It is easy to verify that $\mathbb{E}[Z]$ is dominated by the terms that correspond to two paths of length t which overlap on a single edge. Applying Markov's inequality, see Theorem 1, we obtain

$$\Pr[Z \geq \mathbb{E}[Z] \log n] \leq \frac{1}{\log n} = o(1),$$

This suggests setting $\gamma = n^{t-3/2} \sqrt{\log n}$ since then $|A_{uv}^t - \bar{A}_{uv}^t| \leq n^{t-3/2} \sqrt{\log n}$ whp. The gap requirement is

$$\frac{n^{t-1}(p-q)^t}{k^{t-1}} > 4n^{t-3/2} \sqrt{\log n} \rightarrow p - q = \Omega\left(\left(\sqrt{\frac{\log n}{n}}\right)^{1/t}\right).$$

This proves our claim, as for $t = 2$ we obtain the best possible gap.

4 Experimental Results

In this Section we present our experimental findings. We refer to our method as EGYPT (*Efficient Graph Partitioning*). Specifically, Section 4.1 describes the experimental setup. Section 4.2 presents the simulation results which verify the value of our proposed method. Finally, Section 4.3 provides an application of our method on real-data.

4.1 Datasets

The experiments were performed on a single machine, with Intel Xeon CPU at 2.83 GHz, 6144KB cache size and 50GB of main memory. We have implemented EGYPT in both JAVA JDK 1.6 and MATLAB R2011a. The

method we use to compare against for the 1-path classification is LWD, c.f. [SK12], the single streaming graph partitioning method for which we have theoretical insights [Sta14]. The simulation results in Section 4.2 were obtained using the JAVA code. The results in Section 4.3 were obtained from the MATLAB implementation.

Synthetic data: We generate random graphs according to the planted partition model. We fix $q = 0.05$ and we range the gap from 0.05 until 0.95 with a step of 0.05. This results in pairs of (p, q) values with p ranging from 0.1 until 1. and k ranging from 2 to 16 as successive powers of 2 are qualitatively identical. Notice that given that $p, q = \Theta(1)$ despite the small number of vertices the number of edges is large ranging from 2 to 9.2 million edges.

Parameter B in the interval $\{50, 100, 200, \dots, 1000\}$. The imbalance tolerance ρ was set to 1, demanding equally sized clusters (modulo the remainder of n divided by k). Since the function Ψ is known, i.e., groundtruth is available, we measure the precision of the algorithm, as the percentage of the $\binom{n}{2}$ relationships that it guesses correctly. The success of the algorithm is also judged in terms of the fraction of edges cut λ .

Real data: An interesting question is what kind of real-world graphs does the planted partition model capture? Real-world networks are not modelled by the planted partition model [FGT13]. However, nearest neighbor graphs of well-clustered data points appear to be modelled by the planted partition model we analyzed in Section 3. As a cloud of points, we use a perfectly balanced set of 50 000 digits (5 000 digits for each digit 0, 1, ..., 9) from the MNIST database [yan]. Each digit is a 28×28 matrix which is converted in a 1-dimensional vector with 784 coordinates.

4.2 Simulations

Figures 1(a) and (b) plot the average fraction of edges cut λ and the average precision versus the gap $p - q$ for LWD and three runs of EGYPT with different seed sizes, $B \in \{100, 500, 1000\}$. We observe that even for a small value of B , the improvement over LWD is significant. Furthermore, we observe that even for $p = 1, q = 0.05$ which corresponds to the 0.95 gap, LWD is not able to output a good quality partition. This shows that the analysis of [Sta14] needs even larger values of n to recover the partition *whp*. Furthermore, as B increases from 100 to 1000 the quality of the final partition improves. It is worth emphasizing that the results we obtain are strongly

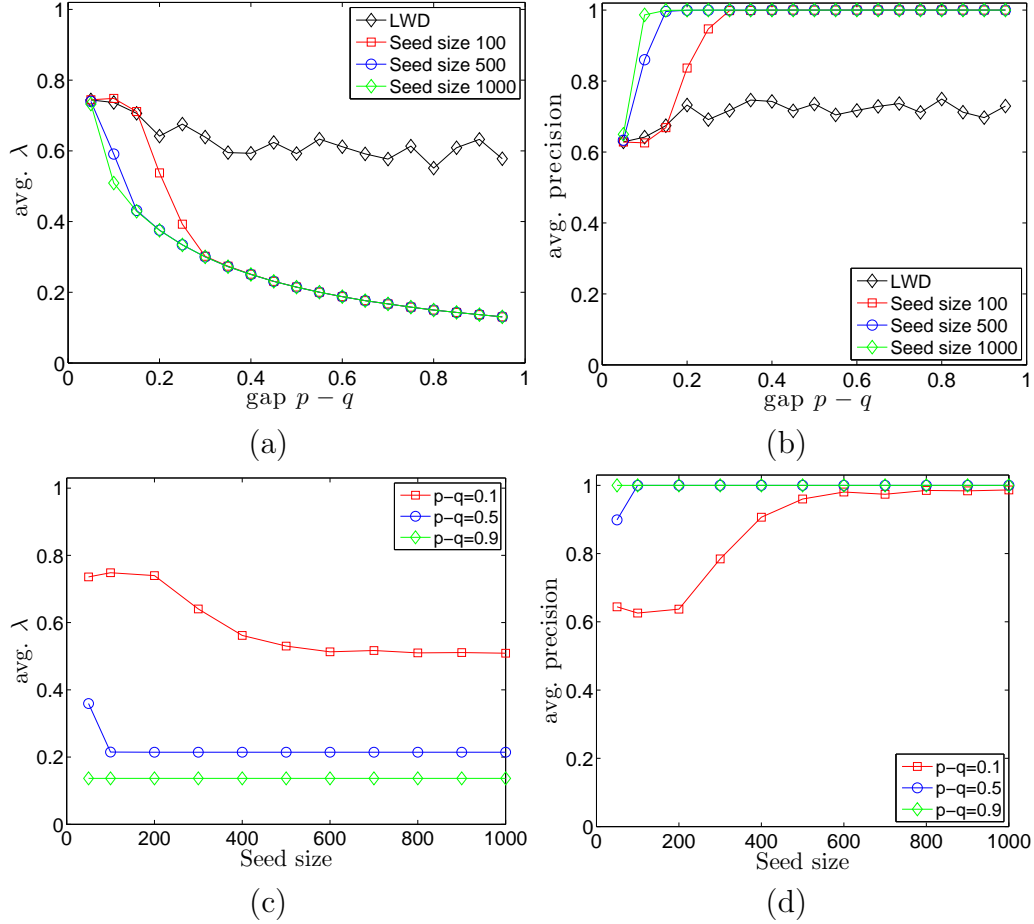


Figure 1: (a) Average fraction of edges cut λ and (b) average precision versus gap $p - q$ for LWD and EGYPT for three different B values (seed sizes). (c) Average precision and (d) average λ , versus B for three different gaps. All data points are averages over five experiments. Observed values are strongly concentrated around their corresponding averages.

concentrated around their corresponding averages. The ratio of the variance over the mean squared was at most 0.0129 and typically of the order 10^{-3} indicating a strong concentration according to Chebyshev's inequality, see Theorem 2.

Finally, Figures 1(c) and (d) plot the average fraction of edges cut λ and the average precision versus the seed size B for three different gaps, averaged over five experiments. Again, data points are concentrated around averages.

As expected, the smaller the gap $p - q$ the larger the parameter B has to be in order to obtain a given amount of precision. When the gap is large even for a small seed size $B = 50$, EGyPT obtains the correct partition. Notice that LWD cannot achieve this level of accuracy. At the same time the computational overhead is negligible, at the order of seconds.




Figure 2: Digits classification

4.3 A Machine Learning Application

We consider a stream of data points where each data point $x \in \mathbb{R}^{784}$ represents a digit from 0 to 9. Whenever a new data point x arrives, we find its $k' = 5$ nearest neighbors among the B first points. This is a variation of the well known k' -nearest neighbor graph [MHvL09]. It is worth outlining that the planted partition model captures some aspects of the resulting k' -nearest neighbor graph. Specifically, the key property we are interested in is the community structure of the graph. For a given set of vertices $U \subseteq V$ the conductance of S is defined $\phi(U)$ is $E(U, \bar{U})/\text{vol}(U)$ where $E(U, \bar{U}) = |\{(u, v) : u \in U, v \in \bar{U}, w(u, v) = 1\}|$ and $\text{vol}(U) = \sum_{v \in U} \deg(v)$. The conductances of the 10 subsets of vertices for each possible digit are 0.0088, 0.0160, 0.0203, 0.0282, 0.0214, 0.0253, 0.0122, 0.0229, 0.0291, 0.0328 respectively. Figure 2 shows the improvement in the precision of the clustering as parameter B increases from 0.1% (50) to 5% (2 500) of the total number

of data points (50 000). The blue straight line shows the performance of LWD. We observe the improvement over LWD even for $B = 50 (= 0.1\% \times 50\,000)$ and the monotone increasing behavior of the precision as a function of B .

5 Conclusions

In this work we introduce the natural idea of higher length walks for streaming graph partitioning, a recent line of research [SK12] that has already had a significant impact on various graph processing systems. We analyze our proposed algorithm EGYPT in the planted partition model.

In future work, we plan to perform an average case performance analysis of EGYPT using a random graph model with power law degree distribution and small separators, e.g., [FFV06, FT14], and to explore the performance of EGYPT on social networks.

6 Acknowledgements

The author would like to thank Alan Frieze and Moez Draief for their feedback on the manuscript.

References

- [bic] “large scale machine learning and other animals. [1](#)
- [BMS⁺] A. Buluc, H. Meyerhenke, I. Safro, P. Sanders, and C. Schulz. Recent advances in graph partitioning. Arxiv 1311.3144. [2.1](#)
- [Bol01] B. Bollobás. *Random graphs*. Cambridge University Press, 2001. [3.2](#)
- [CK01] Anne Condon and Richard M Karp. Algorithms for graph partitioning on the planted partition model. *Random Structures and Algorithms (RSA)*, 18(2):116–140, 2001. ([document](#)), [2.1](#), [3.2](#)
- [CSC⁺13] Rong Chen, Jiaxin Shi, Yanzhe Chen, Haibing Guan, and Haibo Chen. Powerlyra: Differentiated graph computation and partitioning on skewed graphs. 2013. [1](#)

- [DP09] Devdatt P Dubhashi and Alessandro Panconesi. *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press, 2009. 2
- [FFV06] Abraham D Flaxman, Alan M Frieze, and Juan Vera. A geometric preferential attachment model of networks. *Internet Mathematics*, 3(2):187–205, 2006. 5
- [FGT13] Alan M. Frieze, Aristides Gionis, and Charalampos E. Tsourakakis. Algorithmic techniques for modeling and mining large graphs (amazing). In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, page 1523, 2013. 4.1
- [FK02] U. Feige and R. Krauthgamer. A polylogarithmic approximation of the minimum bisection. *SIAM J. Comput.*, 31(4):1090–1118, April 2002. 2.1
- [FT14] Alan Frieze and Charalampos E Tsourakakis. Some properties of random apollonian networks. In *Internet Mathematics*. 2014. 5
- [GLG⁺12] J.E. Gonzalez, Y. Low, H. Gu, D. Bickson, and C. Guestrin. Powergraph: Distributed graph-parallel computation on natural graphs. In *Operating Systems Design and Implementation (OSDI)*, 2012. 1
- [goo] <http://googleblog.blogspot.co.uk/2008/07/we-knew-web-was-big.html>. 1
- [Jer98] G. Jerrum, M. and Sorkin. The metropolis algorithm for graph bisection. *Discrete Applied Mathematics*, 82(1):155–175, 1998. 2.1
- [KL70] B. W. Kernighan and S. Lin. An efficient heuristic procedure for partitioning graphs. *Bell Syst. Tech. J.*, 49(2):291–307, February 1970. 2.1
- [KNS09] R. Krauthgamer, J. (S.) Naor, and R. Schwartz. Partitioning graphs into balanced components. In *Symposium on Discrete Algorithms (SODA)*, pages 942–949, 2009. 2.1

- [Knu07] Donald E Knuth. Seminumerical algorithms. 2007. [4](#)
- [KR04] A. Konstantin and H. Räcke. Balanced graph partitioning. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 120–124, 2004. [2.1](#)
- [KTF09] U Kang, C. E. Tsourakakis, and C. Faloutsos. Pegasus: A petascale graph mining system. In *IEEE International Conference on Data Mining (ICDM)*, 2009. [1](#)
- [LGK⁺10] Y. Low, J. Gonzalez, A. Kyrola, D. Bickson, C. Guestrin, and J. M. Hellerstein. Graphlab: A new framework for parallel machine learning. In *Proceedings of the 2009 ACM SIGMOD International Conference on Management of data*, pages 340–349, 2010. [1](#)
- [LS90] László Lovász and Miklós Simonovits. The mixing rate of markov chains, an isoperimetric inequality, and computing the volume. In *Annual Symposium on Foundations of Computer Science (FOCS)*, pages 346–354. IEEE, 1990. [1](#), [2.1](#)
- [MAB⁺10] G. Malewicz, M. H. Austern, A.J.C Bik, J.C. Dehnert, I. Horn, N. Leiser, and G. Czajkowski. Pregel: a system for large-scale graph processing. In *SIGMOD '10*, pages 135–146, 2010. [1](#)
- [McS01] Frank McSherry. Spectral partitioning of random graphs. In *Annual Symposium on Foundations of Computer Science (FOCS)*, pages 529–537. IEEE, 2001. [\(document\)](#), [2.1](#), [3.2](#)
- [MHvL09] M. Maier, M. Hein, and U. von Luxburg. Optimal construction of k-nearest-neighbor graphs for identifying noisy clusters. *Theoretical Computer Science*, 410(19):1749–1764, 2009. [4.3](#)
- [MU05] M. Mitzenmacher and E. Upfal. *Probability and computing: Randomized algorithms and probabilistic analysis*. Cambridge University Press, 2005. [2](#)
- [NU13] J. Nishimura and J. Ugander. Restreaming graph partitioning: Simple versatile algorithms for advanced balancing. In *ACM KDD*, 2013. [2.1](#)

- [SK12] I. Stanton and G. Kliot. Streaming graph partitioning for large distributed graphs. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2012. ([document](#)), 1, 2.1, 4.1, 5
- [SKK00] K. Schloegel, G. Karypis, and V. Kumar. Parallel multilevel algorithms for multi-constraint graph partitioning (distinguished paper). In *Proceedings from the 6th International Euro-Par Conference on Parallel Processing*, Euro-Par '00, pages 296–310, 2000. 2.1
- [ST08] Daniel A Spielman and Shang-Hua Teng. A local clustering algorithm for massive graphs and its application to nearly-linear time graph partitioning. *arXiv preprint arXiv:0809.3232*, 2008. 1, 2.1
- [Sta14] Isabelle Stanton. Streaming balanced graph partitioning algorithms for random graphs. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1287–1301, 2014. ([document](#)), 1, 2.1, 3.2, 4.1, 4.2
- [TBG⁺13] Charalampos Tsourakakis, Francesco Bonchi, Aristides Gionis, Francesco Gullo, and Maria Tsiarli. Denser than the densest subgraph: extracting optimal quasi-cliques with quality guarantees. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 104–112. ACM, 2013. 2.1
- [TGRV14] C. E. Tsourakakis, C. Gkantsidis, B. Radunovic, and M. Vojnovic. Fennel: Streaming graph partitioning for massive distributed graphs. *Seventh ACM International Conference on Web Search and Data Mining, WSDM 2014*, 2014. ([document](#)), 2.1
- [Vu14] Van Vu. A simple svd algorithm for finding hidden partitions. *CoRR*, abs/1404.3918, 2014. 2.1
- [yan] <http://yann.lecun.com/exdb/mnist/>. 4.1
- [ZW04] Hanson Zhou and David Woodruff. Clustering via matrix powering. In *Symposium on Principles of Database Systems (PODS)*, 2004. 1, 2.1, 3.2, 3.5, 6